

# On the Continuous CNN Problem

John Augustine\* and Nick Gravin

School of Physical and Mathematical Sciences  
Nanyang Technological University  
Singapore 637371.  
jea@ics.uci.edu, ngravin@gmail.com

**Abstract.** In the (discrete) CNN problem, online requests appear as points in  $\mathbb{R}^2$ . Each request must be served before the next one is revealed. We have a server that can serve a request simply by aligning either its  $x$  or  $y$  coordinate with the request. The goal of the online algorithm is to minimize the total  $L_1$  distance traveled by the server to serve all the requests. The best known competitive ratio for the discrete version is 879 (due to Sitters and Stougie).

We study the continuous version, in which, the request can move continuously in  $\mathbb{R}^2$  and the server must continuously serve the request. A simple adversarial argument shows that the lower bound on the competitive ratio of any online algorithm for the continuous CNN problem is 3. Our main contribution is an online algorithm with competitive ratio  $3 + 2\sqrt{3} \approx 6.464$ . Our analysis is tight. The continuous version generalizes the discrete orthogonal CNN problem, in which every request must be  $x$  or  $y$  aligned with the previous request. Therefore, Our result improves upon the previous best competitive ratio of 9 (due to Iwama and Yonezawa).

## 1 Introduction

The  $k$ -server problem has been influential in the development of online algorithms [3]. We have  $k$  servers that can move around a metric space. Requests arrive in an online manner on various locations in the metric space. After each request arrives, one of the  $k$  servers must move to the request location. The online algorithm must make this decision without any knowledge of the future requests. The objective is to minimize the sum of the distances traveled by the  $k$  servers.


A natural variant of the  $k$ -server problem, the (discrete) CNN problem, was introduced by Koutsoupias and Taylor [4]. The name derives from the following illustrative example: consider a sequence of newsworthy events that occur in street intersections in Manhattan. A CNN news crew must cover these events with minimal movement. Since they have powerful zoom lenses, they only need to be at some point on either one of the two cross streets. More formally, we are

---

\* Work done in part while at Tata Research Development and Design Centre, Pune, India.

given a sequence of requests as points from  $\mathbb{R}^2$  that appear in an online manner. We have one server that can move around in  $\mathbb{R}^2$ . To serve a request, the server must merely align itself to the  $x$  or  $y$  coordinate of the request. The objective is to minimize the total distance traveled by the server in  $L_1$  norm.

There is an equivalent alternative definition that is also used in literature in which, instead of a single server that can move in 2D, we have two independent servers with one restricted to move along the  $x$ -axis, while the other is restricted to move along the  $y$ -axis. Given an online request at  $(a, b)$ , either the  $x$ -axis server must move to  $x = a$  or the  $y$ -axis server must move to  $y = b$ . The objective is to minimize the sum of distances moved by either servers. Notice that the two independent servers in different dimensions are equivalent to a single server that can move around in both dimensions. For this reason, the CNN problem is also called sum of two 1-server problems [4].



**Fig. 1.** Illustration for the two server definition of the continuous CNN problem.

We introduce the continuous version of the CNN problem. We use the alternative two server definition to illustrate the continuous version. Consider the problem of covering the activities of a soccer match (see Figure 1). For the sake of simplicity in our illustration, let us have two cameras on rails, one along the length (i.e., the  $x$ -axis server) and the other along the breadth (the  $y$ -axis server) of the field. Their orientations are fixed perpendicular to the direction of movement (of course, pointing into the field). As the ball is kicked around, at least one of the two cameras must track the ball continuously. Informally, the input is a request point moving along a continuous trajectory that is revealed in an online manner and a server must *continuously* align itself to the  $x$  or  $y$  coordinate of the request.

We say that two points are  $x$ -aligned (respectively,  $y$ -aligned) if they share the same  $x$  (respectively,  $y$ ) coordinate. Also, we say that two points are aligned if they are either  $x$ -aligned or  $y$ -aligned. We are now ready to formally define the continuous CNN problem. For this formal definition (and for the rest of the paper) we have a single server that can move around in 2D space. Our input is an online sequence of pairs  $r_i = (p_i, \mathbf{d}_i)$ , where  $p_i$  is a point on  $p_{i-1} + t\mathbf{d}_{i-1}$ ,  $t \geq 0$ , and  $\mathbf{d}_i$  is a unit vector in some arbitrary direction. (In the soccer illustration,  $p_i$  is the point on the previous trajectory where the ball is intercepted and  $\mathbf{d}_i$  is the

new direction in which it is kicked.) Without loss of generality, the first point is assumed to be the origin. The server also starts at the origin. When an input pair  $(p_i, \mathbf{d}_i)$  is revealed, the server and  $p_i$  are already aligned. The online algorithm must then commit to a continuous trajectory  $T_i(t)$  of the server parameterized by  $t$  such that for all  $t \geq 0$ ,  $T_i(t)$  is aligned with  $p_i + t\mathbf{d}_i$ . After the online algorithm commits, the next request  $(p_{i+1}, \mathbf{d}_{i+1})$  arrives, the online server moves to the point on  $T_i$  that aligns with  $p_{i+1}$  along the trajectory  $T_i$ . The objective is to minimize the total distance traveled by the server in  $L_1$  norm.

*History of CNN problems:* The discrete version of the CNN problem was discussed in several conferences and seminars in the late 1990s without any breakthroughs [1, 2]. It was formally introduced by Koutsoupias and Taylor [4]<sup>1</sup>. They conjectured that this problem has a competitive algorithm along with a lower bound of  $6 + \sqrt{17}$  on the competitive ratio of any deterministic online algorithm. Their conjecture was proved affirmatively in [5] by Sitters, Stougie, and de Paepe, albeit, with an algorithm that was  $10^5$ -competitive. For a fascinating discussion of the prevailing understanding of this problem in 2003, see [1]. Eventually, Sitters and Stougie [6] made further improvements and provided a 879-competitive algorithm. In fact, their work focussed on the *generalized  $k$ -server problem* which can be characterized as the sum of several 1-server problems on arbitrary metric spaces. The orthogonal CNN problem was introduced by Iwama and Yonezawa [2]. Each request (except the first one) must either share the  $x$  coordinate or the  $y$  coordinate with the previous request. With this restriction, they were able to improve the competitive ratio dramatically to 9.

*Our Contribution:* We focus on the continuous CNN problem, which is a generalization of the orthogonal CNN problem. We formalize this in the following Claim (with proof deferred to the Appendix).

**Claim 1.** *Any  $c$ -competitive algorithm  $\mathcal{A}$  for the continuous CNN problem can be applied to the orthogonal CNN problem in a manner that preserves the competitive ratio.*

A typical adversarial argument gives us the following lower bound on the competitive ratio.

**Claim 2.** *If there is a  $c$ -competitive algorithm for the continuous CNN problem, then  $c \geq 3$  even on a unit square.*


In Section 2, we introduce a simpler problem called the unit CNN problem and prove a lower bound of 3 on its competitive ratio in Theorem 4. The proof of Theorem 4 can be easily adapted for Claim 2.

We now provide an example that informally illustrates how we get a lower bound of 3 on the competitive ratio of the continuous CNN problem; see Figure 2. Consider the unit square with both the optimal offline server and the online server at the top-left corner. In this adversarial example, the request moves to the

---

<sup>1</sup> Conference version appeared in STACS 2000.

bottom-right corner so that the online server is forced to choose between either a clockwise or counter-clockwise direction. Assume, without loss of generality, that it chooses the clockwise direction and moves to the top-right. The offline server, however, makes a single move down to the bottom-left. Suppose now the request moves around repeatedly in the left and bottom edges of the unit square, i.e., it makes a repeated “L” shaped move. Clearly, the offline server is already at a “sweet spot” and therefore stays unmoved. The online server, however, must correct its position and move to the sweet spot to offset its disadvantage. Notice that the online server moved three units of distance while the optimal offline server just needed just one.



**Fig. 2.** The figure on the left shows the request trajectory. The figure on the right shows the trajectory of online and offline servers.

The significant contribution of our paper is an online algorithm for the continuous CNN problem with a competitive ratio of  $3 + 2\sqrt{3} = 6.464$ . In light of Claim 1, our result improves upon the 9-competitive algorithm for the orthogonal CNN problem [2]. Our algorithm alternates between two phases, namely, the *bishop phase* and the *rook phase*. Hence, we call it the Bishop-Rook algorithm or just the BR algorithm. Our analysis using a non-decreasing potential function is non-trivial. Finally, we show that our analysis is tight by constructing input instances for which the competitive ratio is realized.

*Organization of the paper:* In Section 2, we introduce a simplified problem called the unit CNN problem. We prove a lower bound of 3 on the competitive ratio of algorithms for the unit CNN problem, from which, the lower bound of the continuous CNN problem follows quite immediately. In Section 3 we present the BR Algorithm for the continuous CNN problem. We analyze the BR algorithm in Section 4 and show that it has a competitive ratio of  $(3 + 2\sqrt{3}) \approx 6.464$ .

## 2 Preliminaries: Lower Bound

The lower bound for the competitive ratio of the continuous CNN problem (see Claim 2) is 3. As it turns out, this lower bound shows up in a much simpler problem that we call the *unit CNN problem*. We provide the lower bound proof for the unit CNN problem (Theorem 4) which easily encompasses Claim 2. For the sake of completeness, we state the upper bound for the unit CNN problem, but defer all details and proofs to the Appendix.

In the unit CNN problem, a sequence of requests,  $(r_1, r_2, \dots)$ , appear online as points in  $\mathbb{R}^2$ . We have a server that must serve each request  $r_i$  by moving to a location that is aligned vertically or horizontally with the request. Each time the server moves either horizontally or vertically, it has to pay \$1 (regardless of the distance moved). If the move has both a horizontal and a vertical component, it pays \$2. The objective of the server is to minimize the total payment. We also consider the restricted version called the *unit orthogonal CNN problem* in which any pair of consecutive requests must either share an  $x$  or  $y$  coordinate.

We begin with a simple observation that we state without proof because, in essence, it has been noticed before by Iwama and Yonezawa [2]. Define a sequence of moves by any algorithm to be *frugal* if, for each move in the sequence, the payment is never more than the minimum required to serve the current outstanding request. Therefore, by definition, no frugal sequence of moves will include a \$2 move.

**Observation 3.** *Given any sequence of moves that pays \$d to serve the input sequence of requests, there is a frugal sequence that also pays at most \$d to serve all the requests.*

**Theorem 4.** *If there is a  $c$ -competitive algorithm for the unit CNN problem, then  $c \geq 3$  even when the requests arrive in an orthogonal manner.*

*Proof.* For this proof, we limit ourselves to input instances in which the requests can only appear in the vertices of the unit square with the origin at its bottom left. Furthermore, we restrict the requests to appear in an orthogonal manner, i.e., each request (except the first one) shares either an  $x$  or  $y$  coordinate with the previous request.

Given Observation 3, we can assume that the server will not move if it is already aligned to the request. Notice that an adversary can force the online algorithm to move at each step. We keep server one step away from each request  $r_i$  by placing next request  $r_{i+1}$  diagonally opposite to the current position of the server.

Given such a sequence of requests, an offline algorithm splits it into consecutive triples. The offline server moves at most once to service each triple  $r_{i-1}$ ,  $r_i$  and  $r_{i+1}$ . If possible, the offline server moves to  $r_i$  in one step — it can service  $r_{i-1}$ ,  $r_i$  and  $r_{i+1}$  from  $r_i$ . Such a one step move to  $r_i$  is not possible only when the server and  $r_i$  are diagonally opposite each other, so the server must be within 1 hop of both  $r_{i-1}$  and  $r_{i+1}$ . It simply serves  $r_{i-1}$  from the current position and then hops to  $r_{i+1}$  to serve both  $r_i$  and  $r_{i+1}$ .

So, for every three steps of the online algorithm, the offline server requires at most 1 move, thereby completing the proof.  $\square$

**Claim 5.** *There exists a 4-competitive algorithm for the unit CNN problem*

**Claim 6.** *There exists a 3-competitive algorithm for the orthogonal unit CNN problem.*


### 3 The BR Algorithm for the Continuous CNN Problem

We now turn our attention to the main problem that we address in this paper — the continuous CNN problem. Recall that we formally defined the input as an online sequence of pairs  $(p_i, \mathbf{d}_i)$ . Informally, we treat the request as a point starting at the origin and moving to each subsequent  $p_i$  in straight line segments whose direction is given by the vector  $\mathbf{d}_i$ . So we use the term *request trajectory* to refer to the path traversed by the request. The server’s trajectory must stay aligned with request trajectory at all times. In this section, we describe the Bishop-Rook algorithm or just the BR algorithm that alternates between two phases, namely, the Bishop phase and the Rook phase. As the name implies, the server moves diagonally during the Bishop phase. In the Rook phase, we treat the horizontal and vertical components of the server separately, leading to movements that mimic Rooks in Chess. The algorithm switches between the phases when appropriate conditions (described subsequently for each phase) are met.


The key intuition behind the algorithm is the following. Suppose the offline server manages to get to a “sweet spot” from which it can align with the request trajectory with little or no movement. Then, the online server also must home into that spot. Iwama and Yonezawa [2] also exploit this idea. They get closer to a potential sweet spot using “L” shaped moves — hence, one can call it the Knight algorithm. To achieve this homing effect in the BR algorithm, we define an **offset** vector at the end of the bishop phase that, when added to the online server’s position, will point to our candidate sweet spot. In the rook phase, we use the **offset** vector to guide the online server to the sweet spot.

**Bishop Phase:** During the bishop phase, as the name implies, the server moves diagonally making a  $45^\circ$  angle with the axes. Without loss of generality, let the point  $p_i$  be at  $(0, 0)$  and the online server be on the non-negative part of  $y$ -axis at  $(0, h)$ , so  $h \geq 0$ ; see Figure 3. Throughout the bishop phase, the server moves in a manner that maintains  $x$ -alignment with the request trajectory. Notice that this defines the  $x$  component of the server movement. To ensure the diagonal movement of the bishop phase, the server also moves in the  $-y$  direction. For every maximal  $\delta x$  that the server moves in either the  $+x$  or  $-x$  direction, the server simultaneously moves a distance  $|\delta x|$  in the  $-y$  direction. If (and when) the position of server and request trajectory coincide, we terminate the bishop phase and switch to the rook phase. Let  $(s_x, s_y)$  be the coordinates of the point at which they coincide. Then, the **offset** vector  $\mathbf{o} = -s_x \mathbf{x}$ , where  $\mathbf{x}$  is the unit vector in the positive  $x$  direction.

**Fig. 3.** Bishop phase



**Fig. 4.** Rook phase without **offset** update



**Fig. 5.** Rook phase showing **offset** update

**Rook Phase:** At the beginning of the rook phase, positions of server and request trajectory coincide. Without loss of generality we assume that **offset** is in the  $-x$  direction. We maintain two invariants throughout the rook phase. However, in so doing, we are judicious with the  $L_1$  distance traveled by the server.

**$y$ -alignment:** The server and request trajectory are always  $y$ -aligned. This fully defines the movement of the server along the  $y$  direction because the server maintains the same  $y$  coordinate as that of the request.

**$x$  coordinate inequality:** The  $x$  coordinate of the server is always less than or equal to the  $x$  coordinate of the request trajectory. This invariant is more subtle. When the  $x$  coordinate of the request trajectory is strictly greater than that of the server, the server's  $x$  coordinate stays unchanged — this is to ensure that we are judicious with the  $L_1$  distance traveled. When the  $x$  coordinates coincide and the request trajectory is moving in the  $-x$  direction, then the server moves along with the request trajectory.

During the rook phase, the **offset** vector  $\mathbf{o}$  decreases whenever the server moves. The rate of decrease depends on the horizontal and vertical components of the movement. The rate at which  $|\mathbf{o}|$  decreases is given by:

$$|\mathbf{o}| \leftarrow \begin{cases} |\mathbf{o}| - (1 + \sqrt{3})|t| & \text{if server and request move a distance } t \text{ vertically} \\ |\mathbf{o}| & \text{if request moves but server does not} \\ |\mathbf{o}| - t & \text{if server and request move a distance } t \text{ horizontally} \end{cases}$$

When  $|\mathbf{o}|$  reaches 0, we switch to the bishop phase. Fig. 4 depicts the working of the rook phase, but does not show the change in **offset**. Fig. 5 shows how the **offset** shrinks as the phase progresses.

## 4 Analysis of the BR Algorithm

To simplify the analysis, we assume that we are working on an instance of the continuous *orthogonal* CNN problem, i.e, all the direction vectors  $\mathbf{d}_i$  are orthogonal with respect to the axes. This does not affect our analysis because any straight line of arbitrary angle can be approximated by a series of infinitesimally small  $x$  and  $y$  components.

Before we proceed with the analysis, we make a simple observation that allows us to insert artificial points into the input sequence. Suppose we are given a sequence of input requests  $I = ((p_1, \mathbf{d}_1), \dots, (p_i, \mathbf{d}_i), (p_{i+1}, \mathbf{d}_{i+1}), \dots)$ . Consider the sequence  $I' = (p_1, \mathbf{d}_1), \dots, (p_i, \mathbf{d}_i), (p'_i, \mathbf{d}_i), (p_{i+1}, \mathbf{d}_{i+1}), \dots$ , where  $p'_i$  lies on the line segment between  $p_i$  and  $p_{i+1}$ . Then any server trajectory for serving the request sequence  $I$  will also serve  $I'$  and vice versa.

Our analysis uses a potential function that is non-decreasing throughout the execution of the algorithm. We define a cycle to be the combination of a bishop phase and the subsequent rook phase. Recall that at the start of a cycle, the **offset** is 0. We re-orient our view such that the next outstanding request is at the origin and the online server is at  $(0, h)$ , where  $h \geq 0$ . When re-orienting our



view, we ensure that the potential remains unchanged. This is shown formally in Remark 1. The potential function  $\Phi$  is a function of the **offset** and the parameters defined as follows:

$\ell^{opt}$  and  $\ell^{on}$  are the distances traveled by the optimal offline server and the online server, respectively,  
 $p^{opt}$  and  $p^{on}$  are the positions of the optimal offline server and the online server, respectively.

The potential function is given by

$$\Phi = (3 + 2\sqrt{3})\ell^{opt} - 3d(p^{on} + \mathbf{o}, p^{opt}) - \ell^{on} - |\mathbf{o}| + f(|\mathbf{o}|, p^{opt}, p^{on}), \quad (1)$$

where  $d(p, q)$  is the  $L_1$  distance between points  $p$  and  $q$ . To define  $f$ , we first define  $h = p_y^{opt} - p_y^{on}$ , where  $p_y^{opt}$  and  $p_y^{on}$  are the  $y$  coordinates of  $p^{opt}$  and  $p^{on}$ . Now,

$$f(\mathbf{o}, p^{opt}, p^{on}) = \begin{cases} 0 & \text{if } h \leq 0 \\ (6 - 2\sqrt{3})h & \text{if } 0 \leq h \leq |\mathbf{o}| \\ (6 - 2\sqrt{3})|\mathbf{o}| & \text{if } |\mathbf{o}| \leq h \end{cases}$$

**Theorem 7.**  $\Phi$  is non-decreasing throughout the execution of the BR algorithm and this implies a competitive ratio of  $(3 + 2\sqrt{3})$ .

We first provide a series of lemmas that lead to the proof of Theorem 7.

**Lemma 1.** If the online server stays still,  $\Phi$  does not decrease.

*Proof.* Note that  $\mathbf{o}$  remains unchanged when the online server stays still. Also, the optimal server either (i) does not move, (ii) moves horizontally (arbitrary distance) or (iii) moves vertically the same distance that the request moves. In all three cases,  $\Phi$  does not decrease.  $\square$

**Corollary 1.** From Lemma 1, it follows that, in the bishop phase,  $\Phi$  does not decrease when request moves vertically.

We define the offset halfplane to be the halfplane  $x \leq p_x^{on}$ . Naturally, its complement is  $x > p_x^{on}$ . Since  $p_x^{on}$  can change as the online server moves, the offset halfplane also changes accordingly.

**Corollary 2.** From Lemma 1, it follows that, in the rook phase,  $\Phi$  does not decrease when request moves horizontally in the complement of the offset halfplane.

*Remark 1.* At the start of each cycle, the axes of the euclidean plane can be redrawn (orthogonally) without changing  $\Phi$ .

*Proof.* At the start of each cycle, **offset** is 0. Therefore, only the first three terms of Equation 1 are non-zero. Those three terms do not change if the axes are redrawn orthogonal to the previous axes.  $\square$

In the rest of the lemmas, since we can insert new points into the request sequence, we show that  $\Phi$  does not decrease for small  $\epsilon$  moves of the request in the direction specified.

**Lemma 2.** *In the bishop phase,  $\Phi$  does not decrease when the request moves a distance  $\epsilon$  in the horizontal direction.*

*Proof.* We treat this proof in cases based on the behavior of the optimal offline algorithm.

**Case:  $p^{opt}$  is unchanged.** This is only possible if  $p^{opt}$  and request are  $y$ -aligned.  $\ell^{opt}$  is unchanged.  $\ell^{on}$  increased by  $2\epsilon$ .  $|\mathbf{o}|$  has changed by at most  $\epsilon$ .  $f = 0$  because  $h \leq 0$ . If the request moves in the same direction as  $\mathbf{o}$ , then  $|\mathbf{o}|$  decreases by  $\epsilon$ .  $d(p^{on} + \mathbf{o}, p^{opt})$  decreased by  $\epsilon$ . Overall,  $\Phi$  does not decrease (see Fig. 6).

**Case:  $p^{opt}$  moves vertically and aligns with request.** This is a composition of Lemma 1 and the previous case (see Fig. 7).

**Case:  $p_y^{opt} \leq p_y^{on}$  and  $p^{opt}$  and request are  $x$ -aligned for the duration of the move.**  $\ell^{opt}$  increases by  $\epsilon$ . If request moves in the same direction as  $\mathbf{o}$ , then  $|\mathbf{o}|$  decreases by  $\epsilon$  and  $d(p^{on} + \mathbf{o}, p^{opt})$  decreases by  $2\epsilon$ , otherwise,  $|\mathbf{o}|$  increases by  $\epsilon$  and  $d(p^{on} + \mathbf{o}, p^{opt})$  is unchanged.  $\ell^{on}$  increases by  $2\epsilon$ .  $f$  remains at 0. Therefore,  $\Phi$  does not decrease (see Fig. 8).

**Case:  $p_y^{opt} \geq p_y^{on}$  and  $p^{opt}$  and request are  $x$ -aligned for the duration of the move.**  $\ell^{opt}$  increases by  $\epsilon$ . If request moves in the same direction as  $\mathbf{o}$ , then  $|\mathbf{o}|$  decreases by  $\epsilon$  and  $d(p^{on} + \mathbf{o}, p^{opt})$  remains unchanged. Otherwise,  $|\mathbf{o}|$  increases by  $\epsilon$  and  $d(p^{on} + \mathbf{o}, p^{opt})$  increases by  $2\epsilon$ .  $\ell^{on}$  increases by  $2\epsilon$ .  $h$  in  $f$  increased by  $\epsilon$  (see Fig. 9).

The easy case is when  $|\mathbf{o}|$  decreases. We assume that either  $|\mathbf{o}| \geq h$  or  $|\mathbf{o}| \leq h$ . Otherwise, we can insert a request when the change happens. With either option, the change in  $f$  term is positive and since  $|\mathbf{o}|$  decreases, one can work out that  $\Phi$  increases.


When  $|\mathbf{o}|$  increases, the analysis tightens. The  $f$  term increases by  $(6 - 2\sqrt{3})\epsilon$  because  $|\mathbf{o}|$  and  $h$  also increase by  $\epsilon$ . Therefore,  $\Delta\Phi = (3 + 2\sqrt{3})\epsilon - 6\epsilon - 2\epsilon - \epsilon + (6 - 2\sqrt{3})\epsilon = 0$ .

**Case:  $p^{opt}$  and request are  $x$ -aligned for the duration of the move.** In this case, we are not restricting the relative locations of  $p^{opt}$  and  $p^{on}$ . In particular,  $p^{on} \geq p^{opt}$  first, then after some point, the inequality is interchanged. If we insert a request at that point, then, this case breaks into the previous two cases.


□

**Lemma 3.** *In the rook phase,  $\Phi$  does not decrease when request moves horizontally into the offset halfplane.*


*Proof.* As the request moves a distance  $\epsilon$ , the online server goes with it. (So, the request does not enter the offset halfplane, but rather pushes it by a distance  $\epsilon$ .) Therefore,  $|\mathbf{o}|$  decreases by  $\epsilon$ .




**Fig. 6.** Case:  $p^{opt}$  is unchanged



**Fig. 7.** Case:  $p^{opt}$  moves vertically



**Fig. 8.** Case:  $p_y^{opt} \leq p_y^{on}$



**Fig. 9.** Case:  $p_y^{opt} \geq p_y^{on}$

**Case:  $p^{opt}$  stays still.** Clearly,  $p^{opt}$  is  $y$ -aligned with the request.  $\ell^{opt}$  and  $d(p^{on} + o, p^{opt})$  are unchanged, but  $\ell^{on}$  increases by  $\epsilon$ . Since  $p^{opt}$  and  $p^{on}$  are  $y$ -aligned,  $f = 0$ . Recall that  $|o|$  decreases by  $\epsilon$ . Therefore,  $\Phi$  is unchanged (see Fig. 10).

**Case:  $p^{opt}$  makes a vertical move after which,  $p^{opt}$  and request are  $y$ -aligned.**

We can assume that  $p^{opt}$  made the jump first before  $p^{on}$  moved along with the request. From Lemma 1,  $\Phi$  does not decrease when  $p^{opt}$  jumped.  $p^{on}$  moving along with the request is handled by the previous case (see Fig. 11).

**Case:  $p^{opt}$  makes an  $x$ -aligned move.**  $\ell^{opt}$  and  $\ell^{on}$  increase by  $\epsilon$ .  $d(p^{on} + o, p^{opt})$  decreased by  $\epsilon$ . Since  $|o|$  decreased by  $\epsilon$ ,  $f$  decreases at most by  $(6 - 2\sqrt{3})\epsilon$  (see Fig. 12). Therefore,

$$\Delta(\Phi) \geq (3 + 2\sqrt{3})\epsilon + 3\epsilon + \epsilon - \epsilon - (6 - 2\sqrt{3})\epsilon \geq 0. \quad \square$$

**Lemma 4.** *In the rook phase,  $\Phi$  does not decrease when request moves vertically.*

*Proof.* Note that all vertical moves of  $\epsilon$  distance in the rook phase decrease  $|o|$  by  $(1 + \sqrt{3})\epsilon$ .




Fig. 10. Case:  $p^{opt}$  stays still




Fig. 11. Case:  $p^{opt}$  moves vertically




Fig. 12. Case:  $p^{opt}$  makes an  $x$ -aligned move

**Case:  $p_y^{opt} \leq p_y^{on}$  and  $p^{opt}$  is  $x$ -aligned and therefore does not move.**  $\ell^{opt}$  is obviously unchanged, but  $\ell^{on}$  increases by  $\epsilon$ .  $d(p^{on} + o, p^{opt})$  decreased by at least  $(1 + \sqrt{3})\epsilon - \epsilon = \sqrt{3}\epsilon$ . Since  $h \leq 0$ ,  $\Delta(f) = 0$ . Therefore,  $\Delta(\Phi) \geq 3\sqrt{3}\epsilon + \sqrt{3}\epsilon > 0$  (see Fig. 13).

**Case:  $p_y^{opt} \geq p_y^{on}$  and  $p^{opt}$  is  $x$ -aligned, so it does not move.**  $p^{on}$  and request move up by  $\epsilon$ . As in the previous case,  $\ell^{opt}$  remains unchanged, but  $\ell^{on}$  increases by  $\epsilon$ . Also,  $|o|$  decreased by  $(1 + \sqrt{3})\epsilon$ .  $d(p^{on} + o, p^{opt})$  decreased by  $(1 + \sqrt{3})\epsilon + \epsilon = 2\epsilon + \sqrt{3}\epsilon$ . Both  $h$  and  $|o|$  decreased, so  $f$  decreased as well by at most  $(1 + \sqrt{3})(6 - 2\sqrt{3})\epsilon$  (see Fig. 14). Therefore,


$$\Delta(\Phi) \geq 3(2 + \sqrt{3})\epsilon - \epsilon + (1 + \sqrt{3})\epsilon - (6 - 2\sqrt{3})(1 + \sqrt{3})\epsilon \geq 0$$

**Case:  $p^{opt}$  is still, but  $p^{on}$  and request start below  $p^{opt}$ , move up and cross over to above  $p^{opt}$ .** This is simply a composition of the above two cases, so  $\Phi$  does not decrease.


**Case:  $p_y^{opt} \geq p_y^{on}$  and  $p^{opt}$  is  $x$ -aligned, so it does not move.**  $p^{on}$  and request move down by  $\epsilon$ .  $\ell^{opt}$  is unchanged, but  $\ell^{on}$  increases by  $\epsilon$ .  $|o|$  decreased by  $(1 + \sqrt{3})\epsilon$ .  $d(p^{on} + o, p^{opt})$  decreased by  $(1 + \sqrt{3})\epsilon - \epsilon = \sqrt{3}\epsilon$ . While  $h$  increases,  $|o|$  decreased. Therefore,  $f$  might decrease, but at most by  $(1 + \sqrt{3})(6 - 2\sqrt{3})\epsilon$ . Therefore,  $\Delta(\Phi) \geq 3\sqrt{3}\epsilon + \sqrt{3}\epsilon - (1 + \sqrt{3})(6 - 2\sqrt{3})\epsilon = 0$  (see Fig. 15).

**Case:  $p^{opt}$  starts out  $y$ -aligned, but it  $x$ -aligns itself to the request with a horizontal move.** This case can be viewed as the composition of two parts.  $p^{opt}$  moves first and  $\Phi$  does not decrease (by Lemma 1). Then,  $p^{opt}$  stays still, but request and server move up. This is the previous case. Hence,  $\Phi$  does not decrease (see Fig. 16).


**Case:  $p^{opt}$  moves vertically (up or down) and stays  $y$ -aligned.**  $\ell^{opt}$  and  $\ell^{on}$  increase by  $\epsilon$ .  $|\mathbf{o}|$  decreases by  $(1 + \sqrt{3})\epsilon$ , but  $d(p^{on} + \mathbf{o}, p^{opt})$  increases by at most  $(1 + \sqrt{3})\epsilon$ . Finally,  $f$  remains at 0. Therefore,  $\Delta(\Phi) \geq (3 + 2\sqrt{3})\epsilon - 3(1 + \sqrt{3})\epsilon - \epsilon + \epsilon + \sqrt{3}\epsilon = 0$  (see Fig. 17).  $\square$




**Fig. 13.** Case:  $p_y^{opt} \leq p_y^{on}$  and  $p^{opt}$  is  $x$ -aligned




**Fig. 14.** Case:  $p_y^{opt} \geq p_y^{on}$  and  $p^{opt}$  is  $x$ -aligned.  $p^{on}$  moves up



**Fig. 15.** Case:  $p_y^{opt} \geq p_y^{on}$  and  $p^{opt}$  is  $x$ -aligned.  $p^{on}$  moves down



**Fig. 16.** Case:  $p^{opt}$   $x$ -aligns with a horizontal move



**Fig. 17.** Case:  $p^{opt}$  moves vertically

*Proof (of Theorem 7).*  $\Phi$  started at 0 and, from Lemmas 1, 2, 3, 4, and Corollary 2, we know that it only increased. Without loss of generality, we can assume that we terminate at the end of the rook phase, at which point, the  $f$  function will evaluate to 0. If we terminate at some other point in the cycle,  $f$  might be non-zero. For the purpose of analysis, we can perform a simple

trick to bring  $f$  to zero without increasing  $\ell^{opt}$ . In particular, we artificially move the request repeatedly in an “L” shaped manner with  $p^{opt}$  at the corner.  $p^{on}$  will home in on this corner point as well and once it coincides with the corner,  $f$  will become zero without incurring any increase in  $\ell^{opt}$ . Since  $\Phi = (3 + 2\sqrt{3})\ell^{opt} - 3d(p^{on} + \mathbf{o}, p^{opt}) - \ell^{on} - |\mathbf{o}| \geq 0$ , and  $d(p^{on} + \mathbf{o}, p^{opt})$  and  $|\mathbf{o}|$  are non negative,  $(3 + 2\sqrt{3})\ell^{opt} \geq \ell^{on}$ .  $\square$

*Remark 2.* The analysis of our algorithm is tight, i.e. there are infinite sequences of requests for which  $\ell^{on} = (3 + 2\sqrt{3})\ell^{opt}$ .

*Proof.* We provide two different sequence of input. Our first sequence starts with  $p^{opt}$  at the origin.  $p^{on}$  and request are at  $(0, 1)$  and we are at the beginning of the bishop phase. Then request moves to  $(0, 0)$  and then to  $(1, 0)$ .  $p^{on}$  makes a diagonal move to  $(1, 0)$ , but  $p^{opt}$  does not move. So,  $\ell^{on}$  has increased by 2, but  $\ell^{opt}$  stays unchanged. At this moment our algorithm is in the beginning of the rook phase with  $\mathbf{o} = -\mathbf{x}$ , where  $\mathbf{x}$  is the unit vector along  $x$ -axis. Request moves from  $(1, 0)$  to  $(1, \frac{1}{1+\sqrt{3}})$ ;  $p^{opt}$  moves together with it from  $(0, 0)$  to  $(0, \frac{1}{1+\sqrt{3}})$ . According to our algorithm  $p^{on}$  should follow after the request and at the end of its move  $\mathbf{o}$  becomes 0. Thus we get to the bishop phase of our algorithm.  $\ell^{opt}$  and  $\ell^{on}$  have both increased by  $\frac{1}{1+\sqrt{3}}$ . Note that we are back in the situation that we started, so the adversary can repeat this sequence ad infinitum. In each cycle,  $\ell^{opt}$  has increased by  $\frac{1}{1+\sqrt{3}}$  and  $\ell^{on}$  increased by  $2 + \frac{1}{1+\sqrt{3}}$  length in total. Therefore,  $\frac{\ell^{on}}{\ell^{opt}} = (2 + \frac{1}{1+\sqrt{3}})(1 + \sqrt{3}) = 3 + 2\sqrt{3}$ .  $\square$

In the proof of Remark 2, we provide two tight examples (although the second example is deferred to the Appendix) to indicate how  $\Phi$  balances between multiple scenarios. We feel that minor adjustments to  $\Phi$  will not reduce the competitive ratio.

## Acknowledgment

We are grateful to Ning Chen, Edith Elkind, Sachin Lodha, Srinivasan Iyengar, Sasanka Roy and Dilys Thomas for useful discussions and ideas.

## References

1. Marek Chrobak. Sigact news online algorithms column 1. *SIGACT News*, 34(4):68–77, 2003.
2. Kazuo Iwama and Kouki Yonezawa. The orthogonal CNN problem. *Inf. Process. Lett.*, 90(3):115–120, 2004.
3. Elias Koutsoupias. The k-server problem. *Computer Science Review*, 3(2):105–118, 2009.
4. Elias Koutsoupias and David Scot Taylor. The CNN problem and other k-server variants. *Theor. Comput. Sci.*, 324(2-3):347–359, 2004.

5. René Sitters, Leen Stougie, and Willem de Paepe. A competitive algorithm for the general 2-server problem. In *ICALP '03: 29th International Colloquium on Automata, Languages and Programming, Malaga, Spain*, pages 624–636, 2003.
6. René A. Sitters and Leen Stougie. The generalized two-server problem. *J. ACM*, 53(3):437–458, 2006.

## Appendix

*Proof (of Claim 1).* Recall that, in the orthogonal CNN problem, any two consecutive pairs of requests must share either the same  $x$  or  $y$  coordinate. Given an online input sequence  $I_{ortho} = (p_1, p_2, \dots)$ , we must construct an online sequence of requests  $I_{cont}$  for the continuous CNN problem. When a new request  $p_i$  arrives in  $I_{ortho}$ , we construct the next request in  $I_{cont}$  as follows:  $(p_{i-1}, \mathbf{d}_{i-1})$ , where  $\mathbf{d}_{i-1} = \frac{p_i - p_{i-1}}{|p_i - p_{i-1}|}$  is the unit vector in new direction. Clearly, the request trajectories in both the orthogonal and the continuous CNN problem instances are exactly the same. Therefore, any online algorithm  $\mathcal{A}$  for  $I_{cont}$  can also be used for  $I_{ortho}$ .  $\square$

For the purpose of this proof, we restrict requests to appear in an orthogonal manner, i.e., each request (except the first one) shares either the same  $x$  or  $y$  coordinate as the previous request.

*Second tight instance for Remark 2:* The second sequence is as follows. Both  $p^{opt}$  and  $p^{on}$  be at  $(0, 1)$  in the beginning of the bishop phase. The request is at the origin. Then request moves to  $(1, 0)$  and  $p^{on}$  moves to  $(1, 0)$  and  $\ell^{on}$  increases by 2;  $p^{opt}$  moves to  $(1, 1)$ . At this moment our algorithm is in the beginning of the rook phase with  $\mathbf{o} = -\mathbf{x}$ . Request moves from  $(1, 0)$  to  $(1, -\frac{1}{1+\sqrt{3}})$ ; OPT stays still. Server follows after the request until  $\mathbf{o} = 0$ . Thus we get to the bishop phase of the algorithm. Then request moves to  $(1, 1)$ ;  $p^{opt}$  and  $p^{on}$  stay. Then, request moves to  $(-\frac{1}{1+\sqrt{3}}, 1)$  and finally returns back to  $(1, 1)$ ;  $p^{opt}$  stays still;  $p^{on}$  moves to  $(1, 1)$  at a total length  $3(1 + \frac{1}{1+\sqrt{3}})$ . At this point, we are again at the beginning of the bishop phase. Then request moves to (say)  $(1, 0)$  and we are back to the starting situation, i.e., we are in the bishop phase with  $p^{opt}$  and  $p^{on}$  coinciding and are at unit distance away from the request. Hence, this cycle can be repeated ad infinitum. Clearly,  $\frac{\ell^{on}}{\ell^{opt}} = 4(1 + \frac{1}{1+\sqrt{3}}) + 1 = 3 + 2\sqrt{3}$ .

## The Unit CNN Problem

*Proof (of Claim 5).* We now provide a 4-competitive online algorithm. Our algorithm works in cycles. In each cycle, the online algorithm pays at most \$4. Therefore, to prove that the algorithm is 4-competitive, we have to simply show that the offline optimal algorithm must pay \$1 per cycle. The intuition behind the algorithm is as follows. In each cycle, for the offline algorithm to avoid paying, there must be a *sweet spot* in  $\mathbb{R}^2$  such that if the offline server were located there, it would not have to move throughout the cycle. The goal of the online algorithm is to discover such a position (if it exists) and reach it in at most four \$1 steps. The next cycle starts either (i) when the online algorithm establishes that there is no sweet spot (so the offline server has moved), or (ii) when the offline server cannot serve a request from the sweet spot (again, requiring the offline server to move).



Formally, the algorithm works as follows. Assume that we are at the start of a cycle and the first request in that cycle,  $r_1$ , has arrived. We also assume that the offline algorithm has positioned its server in the advantageous sweet spot. The online algorithm pays \$1 (if required) and aligns with the  $x_1$ . We assume that  $r_2$  does not share the same  $x$  coordinate with  $r_1$ . If it does, the server need not move and it can be discarded from the input sequence (for analysis purposes). When  $r_2 = (x_2, y_2)$  arrives, the online server moves to  $(x_1, y_2)$  and serves  $r_2$ . If  $y_1 = y_2$ , then the sweet spot (if it exists) is somewhere on  $y = y_1 = y_2$ ; this is case A. Otherwise, the sweet spot is either  $(x_1, y_2)$  or  $(x_2, y_1)$ ; this is case B.

Suppose we are in case A. Then we assume that  $r_3 = (x_3, y_3)$  does not share the  $y$  coordinate with  $r_2$ ; if it did, the online server need not move and  $r_3$  can be discarded for analysis purposes. The sweet spot must be  $(x_3, y_1)$ . The online algorithm can reach it in one step and stays there till a request that it cannot serve from  $(x_3, y_1)$  arrives. The cycle is over.

Suppose we are in case B. If the third event  $r_3 = (x_3, y_3)$  does not share either an  $x$  or  $y$  coordinate with  $r_1$  or  $r_2$ , then clearly, there cannot be a sweet spot; the cycle is over. Suppose, instead, that  $r_3$  shares the  $x$  coordinate with  $r_2$ ; other subcases can be seen symmetrically. Then, the sweet spot is  $(x_2, y_1)$ . Then, the online algorithm pays \$2 and moves to  $(x_2, y_1)$ . Again it stays there until forced to move when the cycle is over. Claim 5 (stated in Section 2) follows in a straightforward manner. Furthermore, there are instances for which the competitive ratio 4 can be realized, but we defer their description to the full version.  $\square$

*Proof (of Claim 6).* Suppose the sequence of events is guaranteed to be orthogonal, i.e., each request shares a coordinate with the previous request. Then, we provide a very simple and tight algorithm with competitive ratio 3. The algorithm is very simple. The online server does not move unless the event is not visible to it. In that situation, it moves to the last event that it could see. By the orthogonality condition, we know that this algorithm is correct. Our claim that this algorithm is 3-competitive is along the same lines as the previous theorem, only simpler.

We claim that for every consecutive sequence of moves worth \$3 in our algorithm, OPT has to do at least one \$1 move. The proof is similar to that of Claim 5. Consider four consecutive orthogonal requests  $(r_0, r_1, r_2, r_3)$ ; note that adjacent requests must be at distinct locations. At the start of the cycle, both online and offline servers are in some position to serve  $r_0$ . We assume for the sake of contradiction that the offline algorithm has positioned itself so it does not have to move for the next three requests. The online server will move from its current position to  $r_0 \rightarrow r_1 \rightarrow r_2$  to serve  $r_1, r_2$ , and  $r_3$ , respectively. For the offline optimal algorithm to have served this sequence without moving, it must be in a position to see all four requests. The only candidate for the first three requests is  $r_1$ , while the only candidate for the second three requests is  $r_2$ . Since we require adjacent requests to be distinct, this is a contradiction.  $\square$